

PX128

(Google's Proposed Redactions)

Reviewers

loth@, nirmaljayaram@, tlipus@, nitish@, amini@
 TBD : adarvish@, vivekrao@, bbesson@, nwolf@, eujin@

Box 1: Introduction

The gTrade team in display ads is responsible for connecting different Google's advertising platforms, namely, Google Display Ads (GDA fka GDN) and Display & Video 360 (DV360 fka DBM) to ad exchanges both Google Owned & Operated (O&O) such as AdMob and AdX and Third Party Exchanges (3PEs) such as Rubicon, Casale Media etc. Publishers partner with ad exchanges to monetize their inventory.

Exchanges use different auction mechanisms to sell their inventory. While AdX and AdMob run 2nd Price (2P) auctions it is typical for 3PEs to use a variety of auction types. Different auction types require different bidding strategies. For example 2P auctions are incentive compatible since the winner is charged the runner up bid. So buyers can bid their true value without worrying about overpaying. Exchanges see the gap between winning and runner up bids and are tempted to close the gap to increase their yield. Many 3PEs run 1st Price (1P) auctions on some portion of their inventory where the winner is charged their bid. In such actions, buyers should shade their bids ($\text{bid} < \text{value}$). The amount of bid shading depends on the competition. For example a clairvoyant bidder would bid 1 cent above the runner up.

Commented [1]: maybe link to some well known doc (e.g., by Hal Varian)

To add to the complexity many exchanges are untruthful about their auction type since the actual auction used is opaque to buyers. They claim be running 2P auctions hoping that buyers don't shade their bids and charge the winner something between the runner up and winning bid. We call these auctions untruthful 2P auctions. Truthful 2P auctions are a rarity on 3PEs.

Any auction where the clearing price (cost) depends on winning bid is not incentive compatible and needs bid shading. I am the Tech Lead for a team of 3 data scientists and 4 SWEs focused on efficiently bidding in such auctions. I work closely with several teams such as AdX Quality, DV360 Infra, DV360 Opt, and GDA Opt.

Box 2. Project Poirot: Bidding into adversarial auctions on 3PEs

Fixed CPM advertisers in DV360 (\$2B ARR) specify one CPM bid in the frontend even though they buy across multiple exchanges. If all exchanges ran truthful 2P auctions then bidding for such advertisers would be trivial. In reality, bidding is complicated since truthful 2P auction is a rarity on 3PEs. Exchanges different in auction mechanisms and competitive landscape.

Poirot was created to help advertisers manage the nuances of bidding on different exchanges. The objective is to maximize advertiser surplus (value - cost). It is based on the explore and exploit methodology. We run exploration experiments with on a grid of bid shading multipliers to explore competitive landscape. Based on these experiments we fit a model for surplus which is used to estimate the surplus maximizing bid shading multipliers which are served in the exploitation experiment. Before I joined gTrade team the initial version of Poirot ([go/dbm-poirot](#)) was launched on an accelerated timeline. The model for surplus was basic and hard to interpret. I led a

concerted effort to improve performance of Poirot with significant improvements in surplus modeling which lead to better fits, interpretability, and robustness.

Key contributions:

- Proposed [go/poirot model rev](#) and launched [ariane/215784 summary doc](#) a new model for surplus that improved both the interpretability and model fit (lowers residual variance by 37%) without increasing model complexity.
- Proposed [metrics](#) for Poirot and created a script to generate advertiser surplus distribution analysis ([cr/181646006](#)) for evaluating experiments. This work uncovered an important insight that Poirot model predictions become inaccurate as we add advertiser features in the model which led to removal of advertiser id in the next launch.
- 2018 saw 3PEs declaring the auction type on a sizable chunk of their inventory as 1P. I mentored and led loth@ to add action type as feature to Poirot [ariane/259738 summary doc](#). Adding auction type in Poirot was highly non trivial since the response (log of relative surplus) becomes infinity on 1P inventory. So we defined a new response which needed a [new model](#) for surplus. We decided to switch from log quadratic model to linear cubic b-spline model. A log model has the largest errors at the maxima due to errors being multiplicative which is at odds with our use case which is to get maxima. The skewness in response become more prominent as we increased the grid size. This was addressed by switching from quadratic to cubic b-spline regression.
- I worked with DV360 eng, PMs, sales, and PR to communicate the impact on Poirot launches to advertisers. I led loth@ to document the efficacy of Poirot in protecting advertisers ([go/poirot-benefits-to-advertisers](#), [go/poirot_for_target](#)). Poirot saves advertisers \$600 M annually.
- Poirot for AwBid: AwBid ([go/awbid](#)) refers to the GDA traffic that bids on 3PEs (\$0.5B ARR). I proposed ([go/poirot-for-awbid](#)) and led haou@ to design ([go/gtrade-marple-design](#)) and launch [ariane/258064](#).

Commented [2]: there was this other metrics analysis that christophe did. my recommendation is to combine these into one doc on poirot benefit analysis.

Impact:

Changes to Poirot resulted in 13.6% increase in advertiser surplus across all 3PEs saving them an estimated \$72M annually. The spend on 3PEs dropped by a whopping 32%. Due to DV360 advertisers being mostly budget constrained this resulted in a mix shift in spend across exchanges in favour of exchanges running truthful 2P auctions such as AdX increasing Google Profit by 3.4% (\$41M ARR). Poirot for AwBid increased advertiser ROI on 3PE slice by 11%.

Box 3. AdX 1st Price Bidder

AdX and AdMob have been running 2P auctions since inception. While 2P auctions were supposed to simplify buying publishers have been finding ways to game it. For example on apps publishers engage in multi-calls where they call AdX/AdMob multiple times with decreasing floors to fish out the highest bid. Web publishers on the other hand have started using header bidding which allows them to pit AdX against 3PEs in a 1P auction. This put AdX at a huge disadvantage since the clearing price from its 2P auction i.e., AdX runner up bid is compared to the 1P bids from 3PEs. So AdX could lose to a 3PE bid even if AdX highest bid was actually higher. To remove this disadvantage "last look" was implemented which allowed AdX to look at the highest 3PE bid

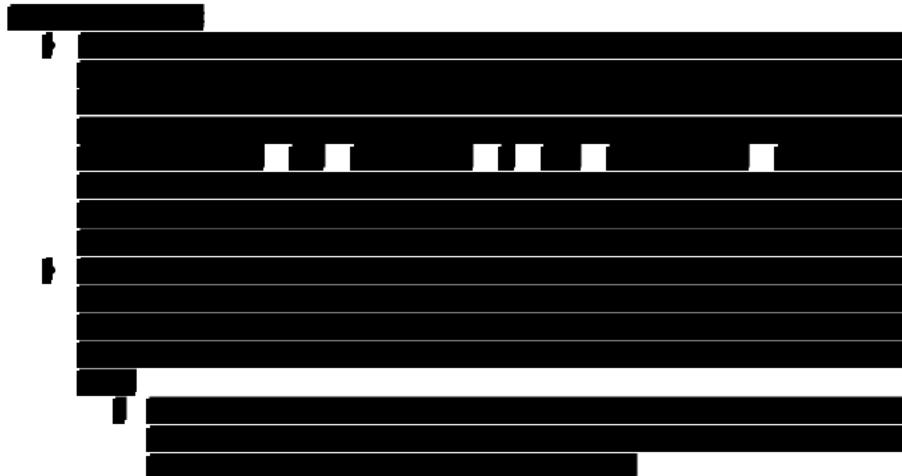
and beat it by 1¢. This was widely perceived in the industry as unfair. To clean up and simplify the sell-side ecosystem AdX and AdMob decided to move to "transparent" 1P (T1P) auction and deprecate last look in Sept 2019.

For Google buyers this has huge ramifications as it impacts more than 50% of their total spend. Unlike in 2P auctions, bidding in 1P is complex as I will explain below. On top of this giving up last look alone leads to revenue losses of 21% and 10% for GDA and DV360 respectively. So we have a lot of ground to make up here.

I led the effort to design, develop and launch ([ariane/311348 go/adx-first-price-bidding-summary](#)) bidders for DV360 and GDA with a team of 3 data scientists and 4 SWEs in a span of 4 quarters. A total of \$5B ARR flows through these bidders. This was a big cross-team effort involving close coordination with DRX quality, DV360 Serving, GDA Opt, and DBM Opt teams. I set the agenda for multiple quarters and drove weekly syncs both within our team and across teams. As we got closer to the launch I ran daily scrum ([notes](#)) to identify and resolve issues quickly.

Key contributions:

Our approach here was fundamentally different from Poirot as T (transparent) in T1P auction means that AdX communicates the Highest Other Bid (HOB) to each buyer post auction. For the winner this is the runner up bid and for the losers this is the winner bid. This information makes exploration experiments unnecessary and allows us to use ML to predict competition and then compute bids in serving at scale and precision which is not possible in Poirot. The DV360 and GDA bidders share this high level methodology but have significant differences which arise from their different business models.



Commented [3]: i see a series of bullets, but no link from bullet to bullet. it's like i did x,y,z

Instead explain it as a story.
To get 1p bid, we need to optimize for adv value. The inputs needed for this are adv value and distribution of competition
This solution is not in closed form so i did..
We build distribution of competition using PEAR

DV3
Once the basic framework was set, we noticed surplus maximizer wasn't giving us the right tradeoffs.
proposed and launched risk averse bidding
Value was providing an important signal for the distribution of competition so we build calibration model

Adwords
On top of the basic framework, there are additional constraints for adwords
(a) incentive compatibility
(b) bernanke to achieve margins

Similar to DV3, we have calibration for predicted competition. In addition, floors can be viewed as comparision as well and some floors are shared with the buyers. This led to floor aware bidding.

Move simulation out as a super bullet and explain this was important for model building

Commented [4]: this was very helpful . I just finished the common part. Do you think its on the right track? too detailed?



DV360 ([go/dv360-1p-bidder-hld](#)): When we ran our first set of experiments with the common framework (surplus maximizing bidding) our experiments showed a healthy surplus but the win rates were very low resulting in significantly loss in volume. Buyers using surplus as utility are risk neutral. Such bidders are willing to gamble on very small chance of winning a very large prize creating St. Petersburg Paradox. I proposed, implemented and launched risk averse bidding of DV360. Risk aversion makes the bidder care more about not losing a query than overpaying slightly. This provided a valuable knob to choose a different tradeoff between surplus and win rate than that offered by surplus maximization.

GDA: On top of the basic framework, GDA has additional constraints on the 1P bidder:

- Preserving incentive compatibility: Most GDA advertisers use auto bidding which assumes incentive compatibility. Hence preserving incentive compatible was critical. I worked with nirmaljayaram@ and tipus@ to come up with a pricing scheme which coupled with surplus maximization bidding makes the mechanism incentive compatible from advertiser point of view. When we proposed this mechanism there was some concern regarding how the surplus extracted was going to be shared between advertisers and Google. It's possible to create an example where surplus maximization + incentive compatible pricing leads to Google keeping all the surplus. I theoretically proved that our mechanism in expectation is incentive aligned with the advertiser ([go/gda-1p-bidder-advertiser-incentive](#)). The only way google extracts a profit is via arbitrage (internal competition is higher than external competition). If we assume no internal competition then all the surplus is passed on to the advertiser and google's margin will be 0 in expectation.
- Buy side margin constraints per publishers: GDA as a business is only allowed to extract an aggregate margin of 15% per publisher. In the 2P world margin was being managed by project Bernanke. The GDA 1p bidder has to solve a constrained optimization problem (surplus max with margin constraint). I guided pooyaj@ to develop Bernanke for 1p ([go/bernanke-for-1p](#)). The elegance of the solution is that it breaks the problem into two

parts which we already know how to solve. First part is to figure out a bidding function that maximizes surplus. Given this function it provides a knob that can be tuned to hit margin via exploration experiments or simulations. The parametrization of the solution for 1P looks very similar to 2P Bernanke but the theoretical arguments that prove its optimality are completely different due to the nonlinear nature of 1p bidding function.

Experiments:

- To allow all buyers (Google and external) to build their 1P bidders, AdX started running 1P auctions on x% of their traffic before launch. Within this x% of 1P traffic, we needed to run many experiments with different bidding strategies. Ensuring that total traffic in our experiments was always x% was going to be impossible since we were going to rapidly iterate. So I designed the externally visible experiment ([go/domains_for_external_expt](#)) using domain which decoupled x% from the total traffic in our bidding experiments.
- Started, debugged and analyzed 50+ experiments to converge on the launch candidate.

Commented [5]: need to explain this better. adx needed to run X% experiments as first price so all buyers get a chance to try out their 1p algorithms. Within this space of X, we need to try out many bidding strategies for GDA and DV. Hence you proposed...

Commented [6]: done

Simulations: Parameter tuning via experiments takes days. So I built a fast and accurate simulation code ([link](#)) which accelerated experiment design and decision making. We were able to tune our experiments on their 2nd day. It also allowed us to learn how to use floors in our bidding ([go/floor_aware_bidding](#)). These simulations showed that using the floors incorrectly can make the mechanism non incentive compatible.

Launch metrics

Compared to second price auction (current production):

- GDA: impression +3.70%, revenue 1.2%, value -5.16%, Google profit neutral.
- DV360: impression +13.43%, revenue 2.33%, NoBc revenue +0.86%, value +1.10%, Google profit -6.33%.

Relative to bid translation (AdX's default first-price-bidding behavior): +11% revenue, +9% net revenue

Commented [7]: rewrite
As part of Adx moving to 1p we are giving up the "last-look advantage". If we use the simple bid translation approach which has been used in the past, we see -30% adwords and -10% dbm revenue. A highly efficient bidder was critical to recover from this handicap. The launched version of the smart bidder, not only recovered this loss, but added..in impressions and revenue.

4. Other Contributions

- Manage 2 DSs.
- 23 DS interviews this cycle
- Mentored nooglers loth@, pooyaj@ and sammolen@.
- Presented @
 - DS summit 2019
 - Ads DS summit Q2 2019 ([slides](#))
 - DV360 Engg ([slides](#))
 - Bidding & Predictions All Hands Q2 2018
- Posters:
 - DS summit 2018